



More Than You Could Want to Know



Vision: Create an e-Infrastructure






The Grid

“Resource sharing & coordinated problem solving in dynamic, multi-institutional virtual organizations”

- ◆ Coordinates distributed resources ...
- ◆ using standard, open, general-purpose protocols and interfaces ...
- ◆ to deliver required qualities of experience



Context

- ◆ Define a service-oriented architecture ...
 - ◇ the key to effective virtualization
- ◆ ... to address vital “Grid” requirements
 - ◇ AKA utility, on-demand, system management, collaborative computing
- ◆ ... building on Web services standards
 - ◇ extending those standards where needed
 - ◇ enabling ubiquity and thus a rich ecology of platform and solution providers
- ➔ The “Open Grid Services Architecture”

Web Services

- ◆ Increasingly popular standards-based framework for accessing network applications
 - ◇ *Service oriented architecture*: everything is a service
 - ◇ Services built in *application server*
- ◆ XML based (eXtensible Markup Language) distributed computing framework
 - ◇ WSDL: Web Services Description Language
 - Interface Definition Language for Web services
 - ◇ SOAP: Simple Object Access Protocol
 - XML-based RPC protocol; common WSDL target
 - ◇ WS-Inspection
 - Conventions for locating service descriptions
 - ◇ UDDI: Universal Desc., Discovery, & Integration
 - Directory for Web services



www.globus.org

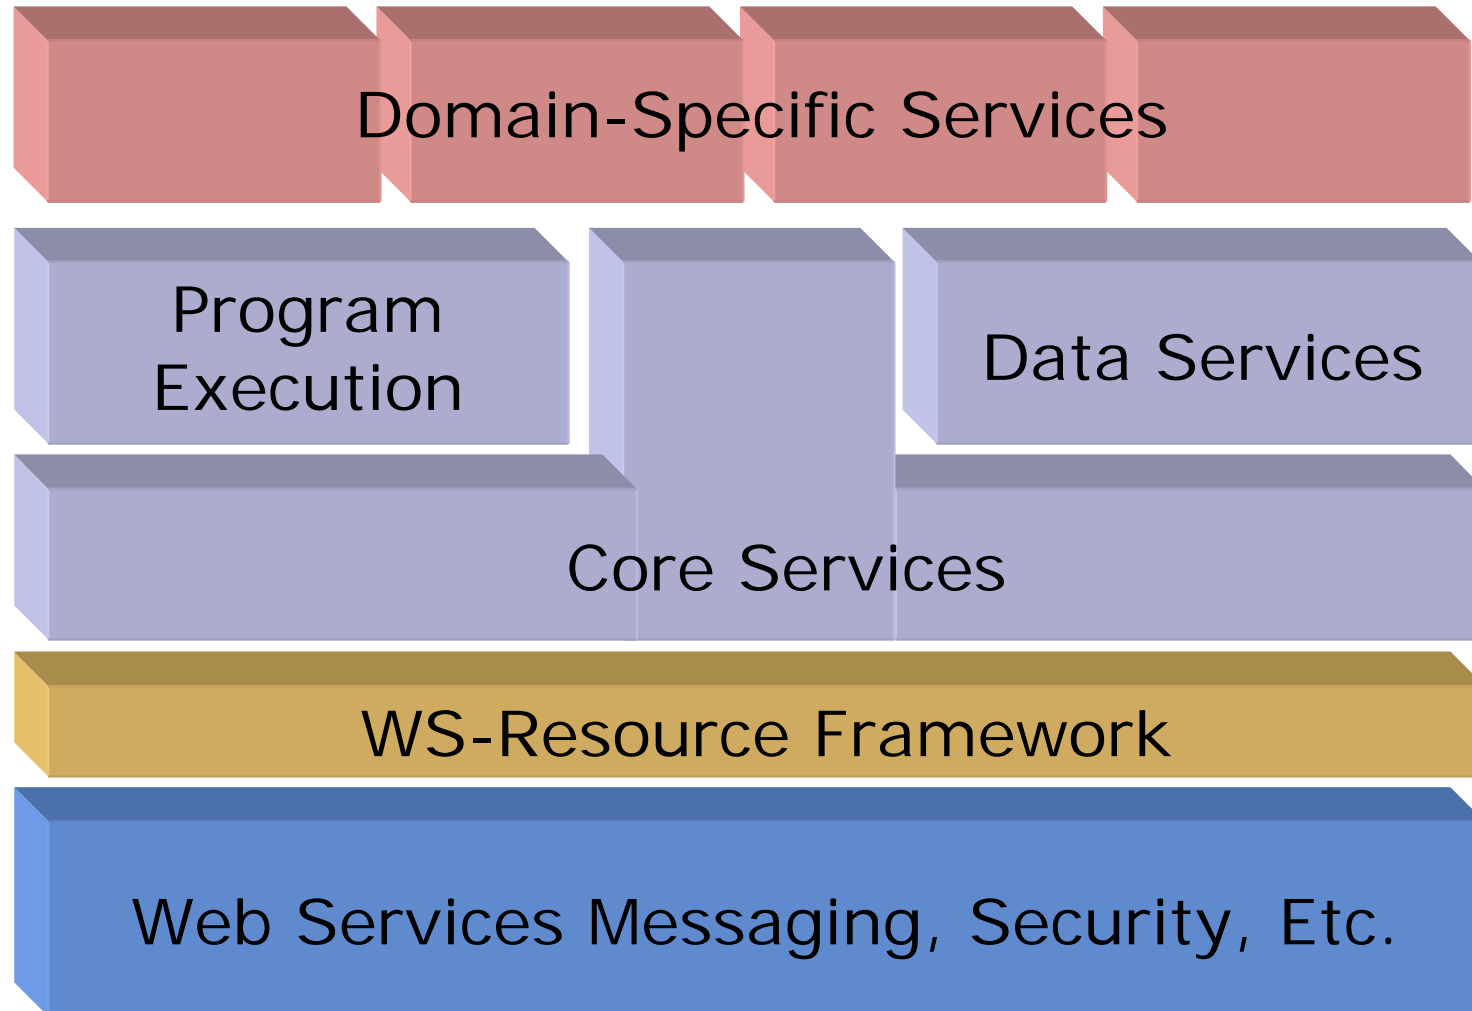
The Grid Enables Transient Services

- ◆ “Web services” address discovery & invocation of persistent services
 - ◇ Interface to persistent state of entire enterprise
- ◆ In Grids, must also support transient services, created/destroyed dynamically
 - ◇ Interfaces to the states of distributed activities
 - ◇ E.g. workflow, video conf., dist. data analysis
- ◆ Significant implications for how services are managed, named, discovered, and used
 - ◇ In fact, much of our work is concerned with the management of services

Open Grid Services Infrastructure (OGSI) to Web Services Reference Framework (WSRF)

- ◆ Globus Toolkit → OGSI → WSRF
- ◆ Announced at GlobusWORLD 2004 (Jan 20-23)
- ◆ The WSRF proposal is a re-factoring of OGSI concepts to align better with Web services
- ◆ Changes from OGSI to WSRF are primarily syntactic
- ◆ Grid adds to Web services in OGSI and WSRF:
 - ◇ The ability to create, address, inspect, discover, and manage stateful resources
 - ◇ In OGSI, these stateful resources are called Grid services
 - ◇ in WSRF, they are called WS-Resources


Grid-Web Services Convergence Completed: A Major Milestone!



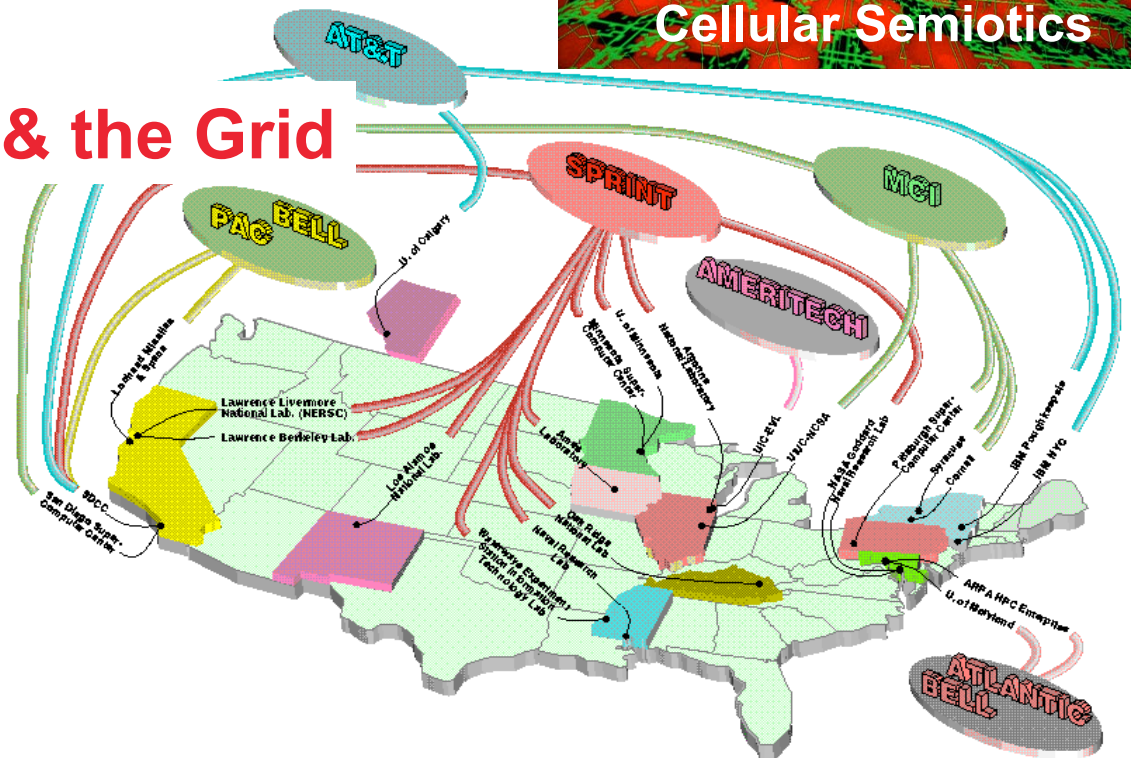
Supercomputing 95

I-Way Featured:

- ◆ Networked Visualization Application Demonstrations
- ◆ OC-3 (155Mbps) Backbone
- ◆ Large-Scale Immersive Displays
- ◆ I-Soft Programming Environment



Led Directly to Globus & the Grid

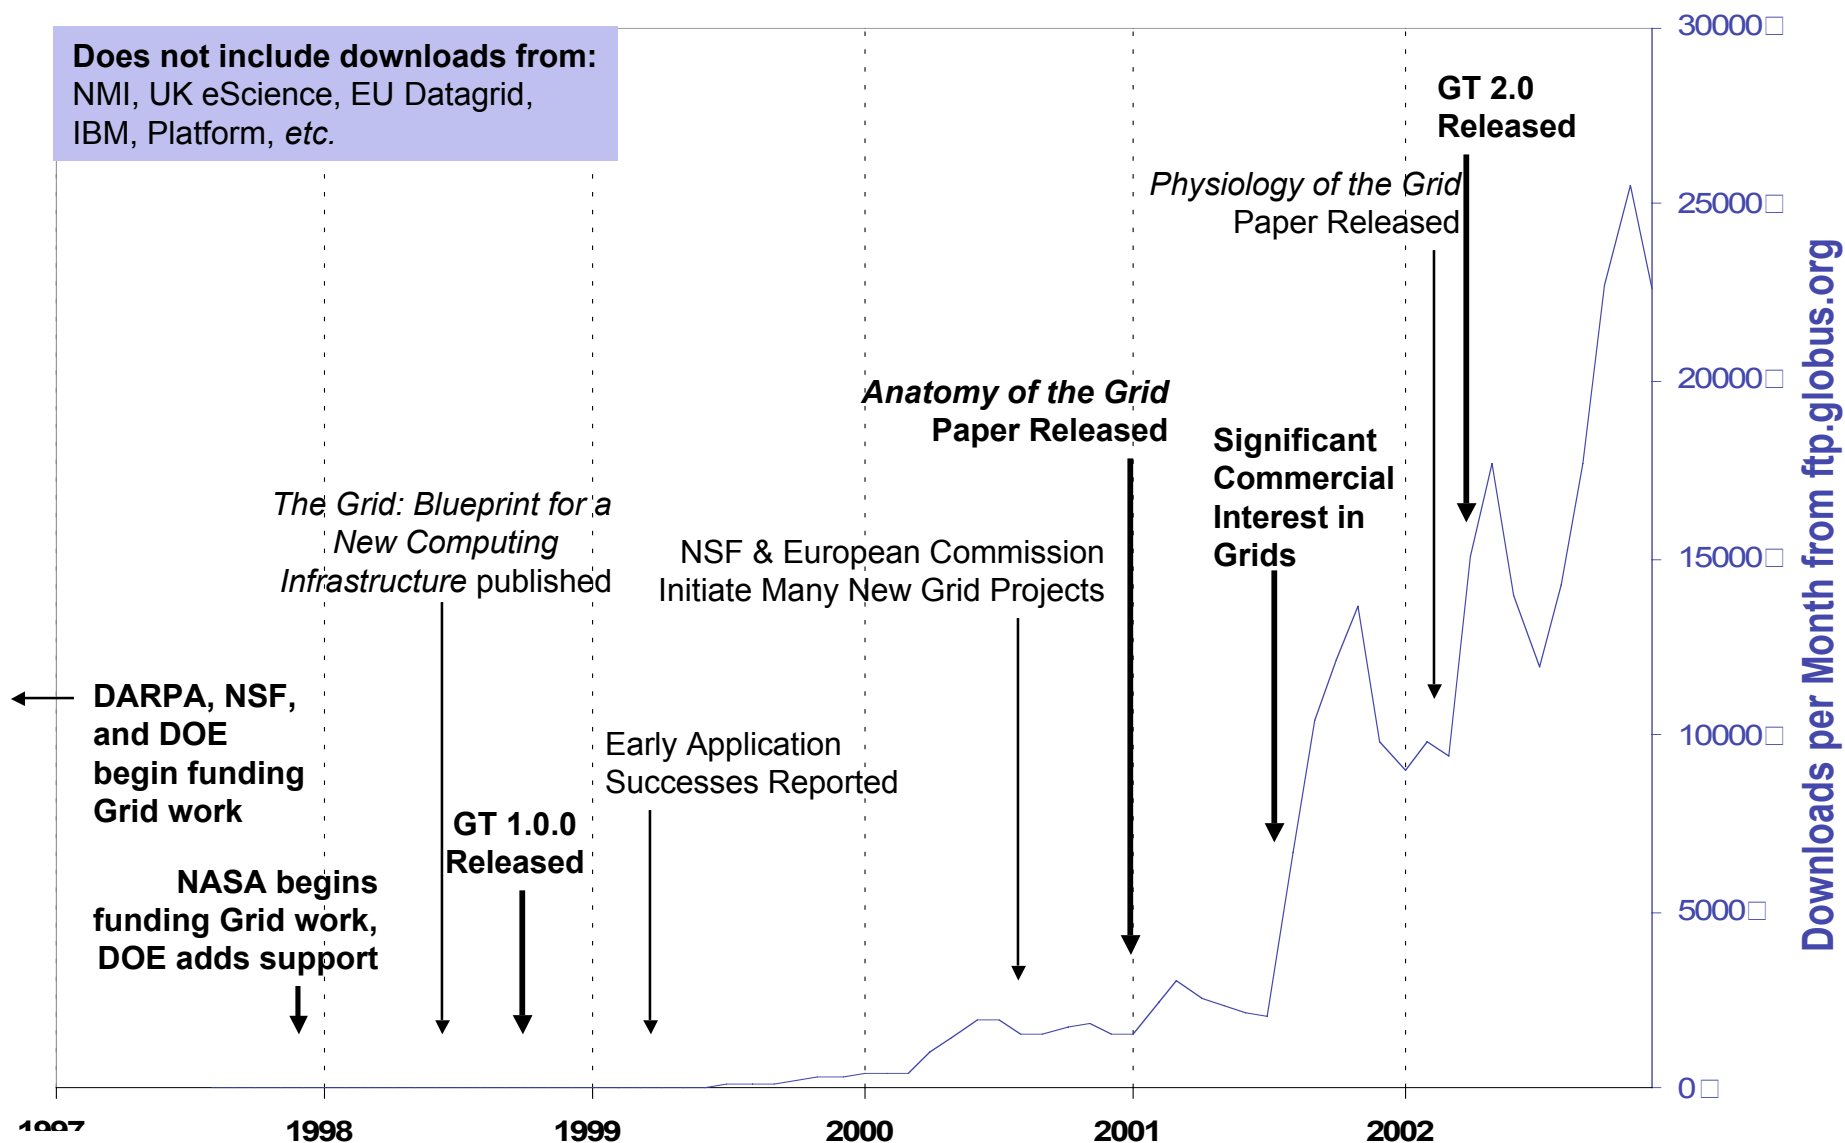


<http://archive.ncsa.uiuc.edu/General/Training/SC95/GII.HPCC.html>



the globus alliance
www.globus.org

Globus Toolkit® History



Partners are Creating Strong GT-Based Grid Solutions



Note that WSRF has little or no effect on solution users!

Charles Schwab

Business Analytics

Challenge

- Reduce the processing time on an existing wealth management application.

Solution

- IBM @server
- Linux
- Globus Toolkit
- IBM Research

"We believe that Grid computing ... has the potential to greatly improve our quality of service and be a truly disruptive technology."

Oren Leiman, Managing Director, Charles Schwab



Technology Benefits:

- Reduced processing time from more than four minutes to fifteen seconds
- Planning to explore leveraging Grid computing into other areas

Business Benefits:

- Potential to increase customer satisfaction by responding to inquiries at a faster pace
- Potential to enable Schwab to provide more robust wealth management applications

© 2003 IBM Corporation

Industry Adopts Grid Technology

I.B.M. Making a Commitment to Next Phase of the Internet

By STEVE LOHR

The New York Times

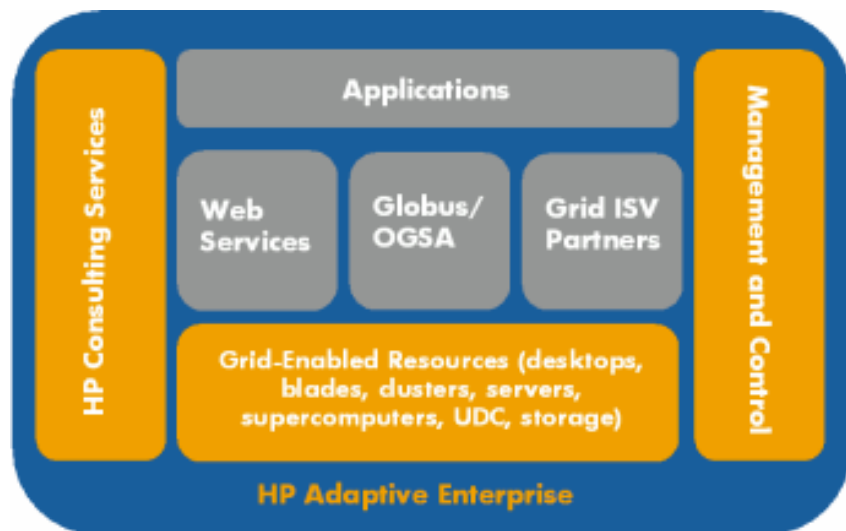
I.B.M. is announcing today a new initiative to support and exploit a technology known as grid computing, which the company and much of the computer research community say is the



Globus Grid Computing—the Next Internet

by John Roy/Steve Milunovich

The Internet was first a network and is now a communications platform. The next evolutionary step could be to a platform for distributed computing. This ability to manage applications and share data over the network is called "grid computing."



ORACLE DATABASE 10^g

THE WORLD'S FIRST

SELF-MANAGING, GRID-READY DATABASE ARRIVES

Oracle's new self-managing database increases performance and availability while enabling commercial grid computing.

There are trends. And there are breakthroughs. In the early to mid 1990s, Oracle foresaw the Internet computing paradigm that organizations of every stripe have now woven into the fabric of their businesses. In the process, IT infrastructure has become extremely critical to the enterprise. "Businesses have become more dependent than ever on their IT systems for everything from day-to-day operations to providing service to their customers and clients," says Satish Kumar, director of product management for Database Management at Oracle. "But many new-generation businesses, such as eBay and Amazon, rely completely on their IT infrastructure's being available—if the system goes away, their entire business is in jeopardy."

In short, says Kumar, IT systems have truly become strategic to the enterprise. And that has had a profound impact on the need for availability, scalability, and high performance. IT systems for organizations of all kinds. Downtime, even for much-needed maintenance, is not an option when a global business must run 24/7.

At the same time, says Kumar, there's growing pressure to maintain profitability amid ever-growing competition in a global economy that continues to tighten its belt. The result, he says, is that "organizations must minimize operating expenses across the board—and IT is no exception," says Kumar.

But as IT systems have become more strategic and integral to the core business, they have also become more complex, more difficult to manage, and more costly. Completely adding to costs across the board, in terms of time, labor, potential failure, and inability to recover from failure effectively. According to Kumar, these are all reasons why "one of the biggest challenges facing most organizations today is managing a strategic part of the business, its IT systems, more effectively than ever—ensuring the highest performance, scalability, and availability—but at a significantly lower cost than before." These are also some of the reasons that commercial grid computing, enabled in part by cost-effective blade servers, is getting so much attention today. The small incremental costs, organizations can gain more processing power to be used by all data center resources, delivering faster performance and high availability and scaling as needed—but only if the software can effectively take advantage of that architecture. Clearly, the time is right for software that monitors and manages itself: software that issues management complexity in a cost-effective manner.

BY KELLI WISETH

The GEON Project

- ◆ Close collaboration between geoscientists and IT to interlink databases and Grid-enable applications
- ◆ “Deep” data modeling of 4D data
 - ◇ Situating 4D data in context—spatial, temporal, topic, process
 - ◇ **Semantic integration** of Geosciences data
 - Logic-based formalisms to represent knowledge and map between ontologies
- ◆ Grid computing
 - ◇ **Workflow-based** specification of computational tasks
 - ◇ Deploying a prototype GEON **Grid**: heterogeneous networks, PoP nodes, Data clusters, and compute clusters
- ◆ Interaction environments
 - ◇ **Information visualization**. Visualization of concept maps
 - ◇ **Remote data visualization** via high-speed networks
 - ◇ Augmented reality in the field



Source: Chaitan Baru, SDSC




CYBERINFRASTRUCTURE FOR THE GEOSCIENCES

A Geoscientist's Information Integration Scenario

*What is the distribution and U/ Pb zircon ages of A-type plutons in VA?
How about their 3-D geometry using gravity data ?
How do the plutons relate to the host rock structures?*

?


Information
Integration




*Digital geologic map
Of Virginia
(plutons in Virginia)*

SiO ₂	72.22
CaO	0.62
K ₂ O	4.88
Ga	21.1
Sr	72.6


*Geochemical
Database
(chemical data)*



*Geophysical Database
(gravity contours)*



*Geochronologic
database
(Concordia)*



*Structure database
(foliation map)*



Virginia Tech & GEON

Source: Chaitan Baru, SDSC



CYBERINFRASTRUCTURE FOR THE GEOSCIENCES

Web Services and GEON

- ◆ GEON software stack is based on OGSII / Globus Toolkit3
- ◆ Each node in the system runs same software stack
 - ◇ Installed via Rocks, and GEON Rolls (a set of GEON-specific RPM's)
- ◆ Approach to Web services
 - ◇ GEON services are Web services with GSI authentication
 - ◇ Will move to WSRF framework when it is mature
 - ◇ Implementing a distributed Web services catalog across GEON
- ◆ Web services management
 - ◇ SDSC currently works with Cal-(IT)² partner, Blue Titan (BT)
 - ◇ BT Network Director provides WS management for services published at webservices.sdsc.edu.
 - ◇ Similar capability being investigated for GEON.
 - ◇ Provides for failover and QoS capabilities



Source: Chaitan Baru, SDSC



CYBERINFRASTRUCTURE FOR THE GEOSCIENCES

See Nov 2003 CACM For Articles on OptIPuter Technologies

The **OptIPuter**: A Revolutionary LambdaGrid Networking Architecture to Support Data-Intensive e-Science Research

Learn about the **OptIPuter** by reading the November 2003 issue
of the *Communications of the ACM* in these articles:

BLUEPRINT FOR THE FUTURE OF HIGH-PERFORMANCE NETWORKING

By MAXINE D. BROWN, GUEST EDITOR

DATA INTEGRATION IN A BANDWIDTH-RICH WORLD

By IAN FOSTER AND ROBERT L. GROSSMAN

TRANSLIGHT A GLOBAL-SCALE LAMBDAGRID FOR E-SCIENCE

By TOM DEFANTI, CÉS-DE LAUT, JOE MAMBRETTI,
KEES NEGGER, AND BILL STAMPAUD

THE OPTIPUTER

By LARRY L. SARRIS, ANDREW A. CHEN, TOM DEFANTI,
LEON TROCH, AND PHILIP M. PAPADOPOULOS

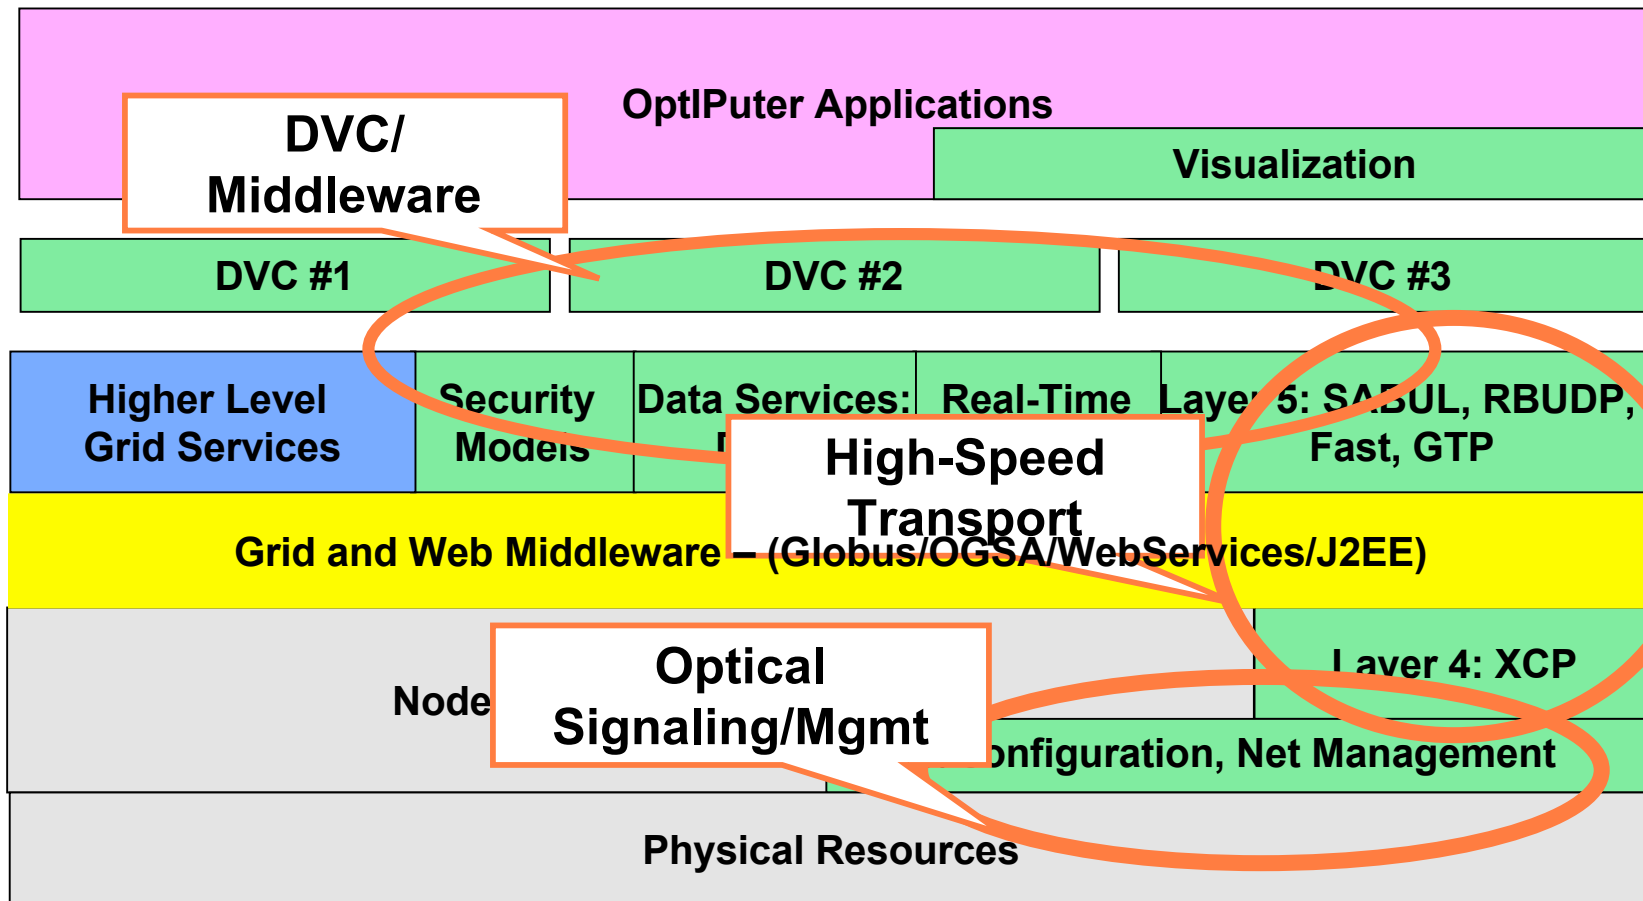
TRANSPORT PROTOCOLS FOR HIGH PERFORMANCE

By AARON FALK, TED FARRIS, JOSEPH RANISTER,
ANDREW CHEN, ROBERT GROSSMAN, AND JASON LEIGH

DATA-INTENSIVE E-SCIENCE FRONTIER RESEARCH

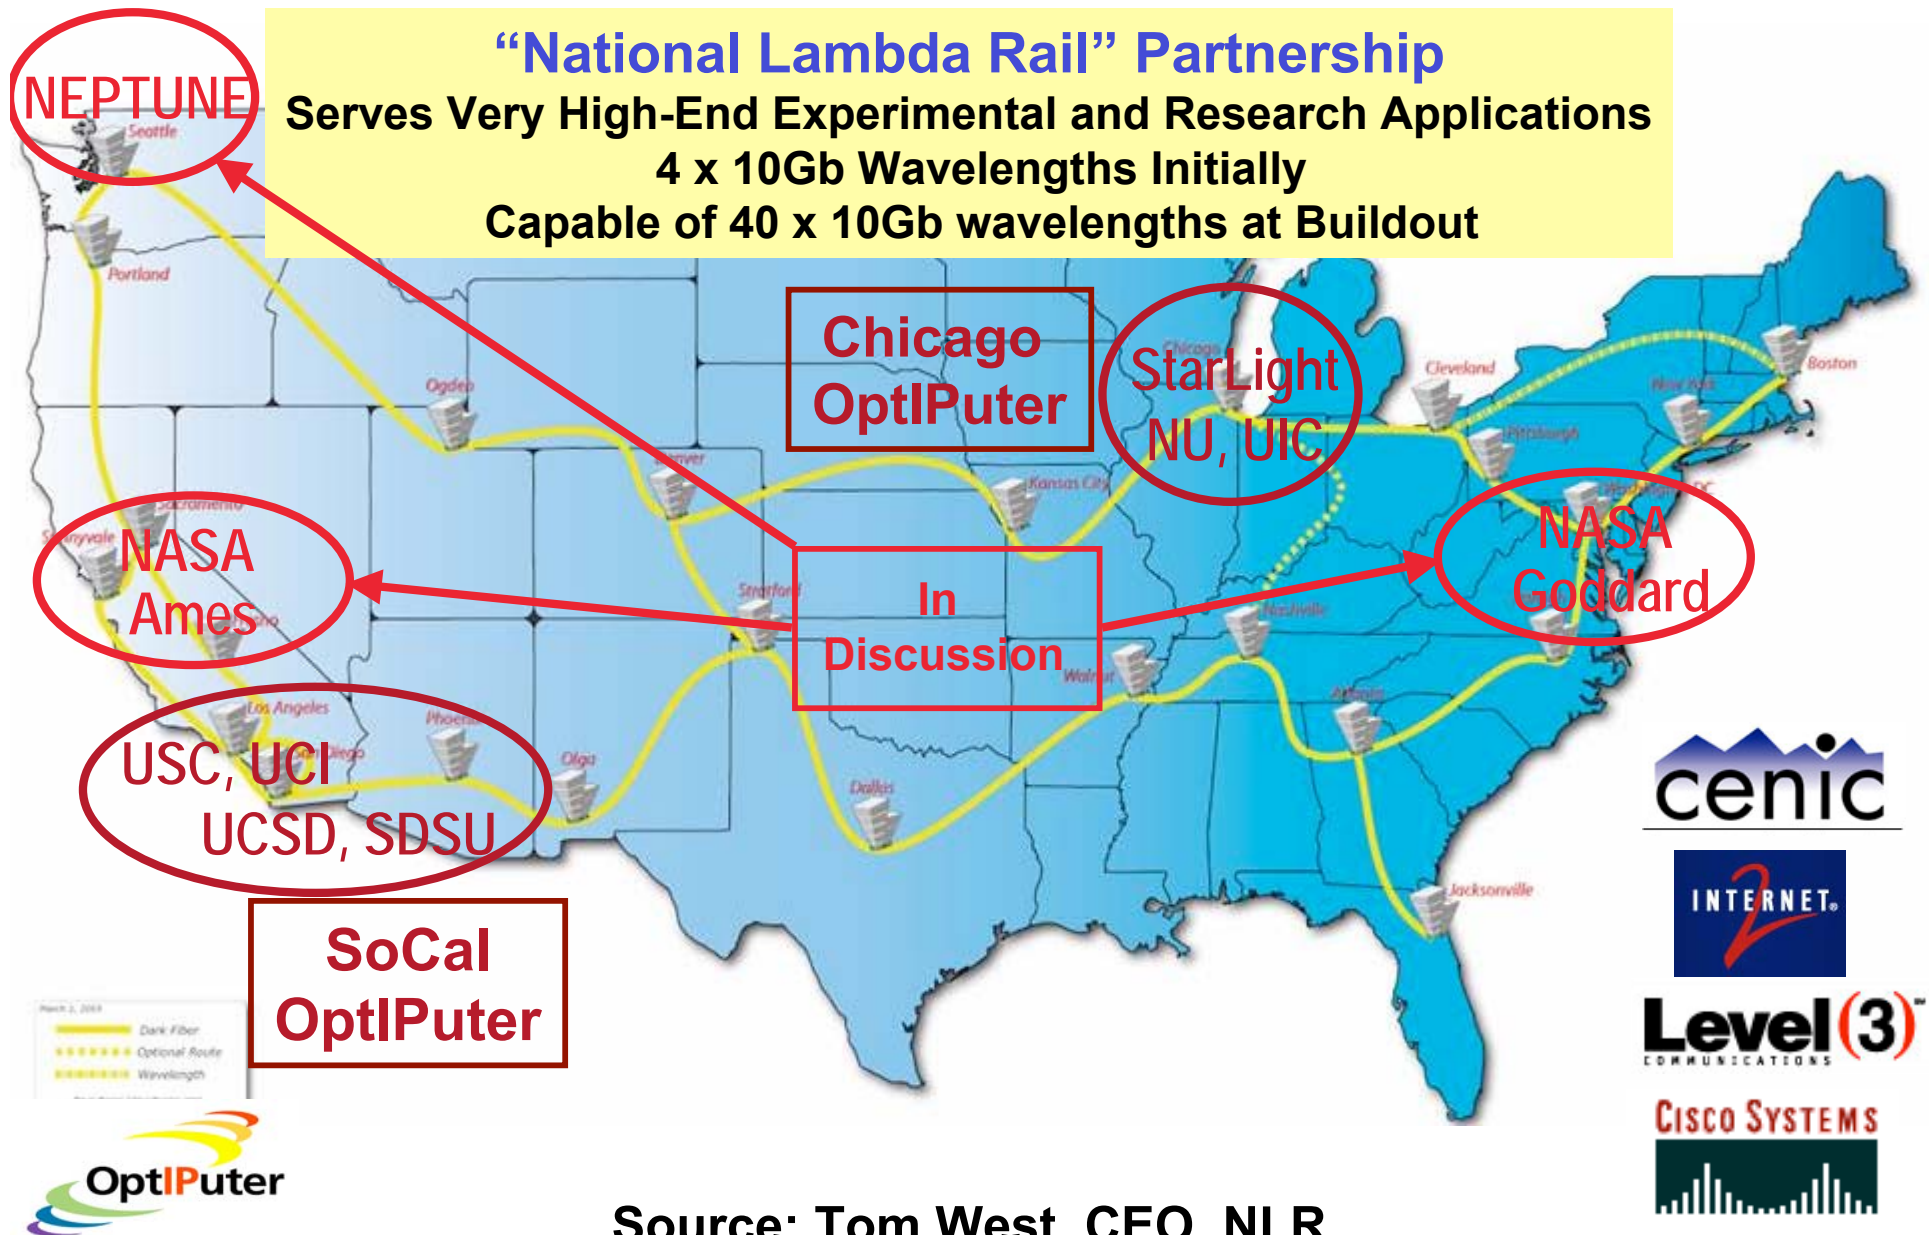
By HARVEY B. NEWMAN, MARK H. BUCKMAN, AND JOHN A. GRIFFIN

OptIPuter Software Architecture for Distributed Virtual Computers v1.1



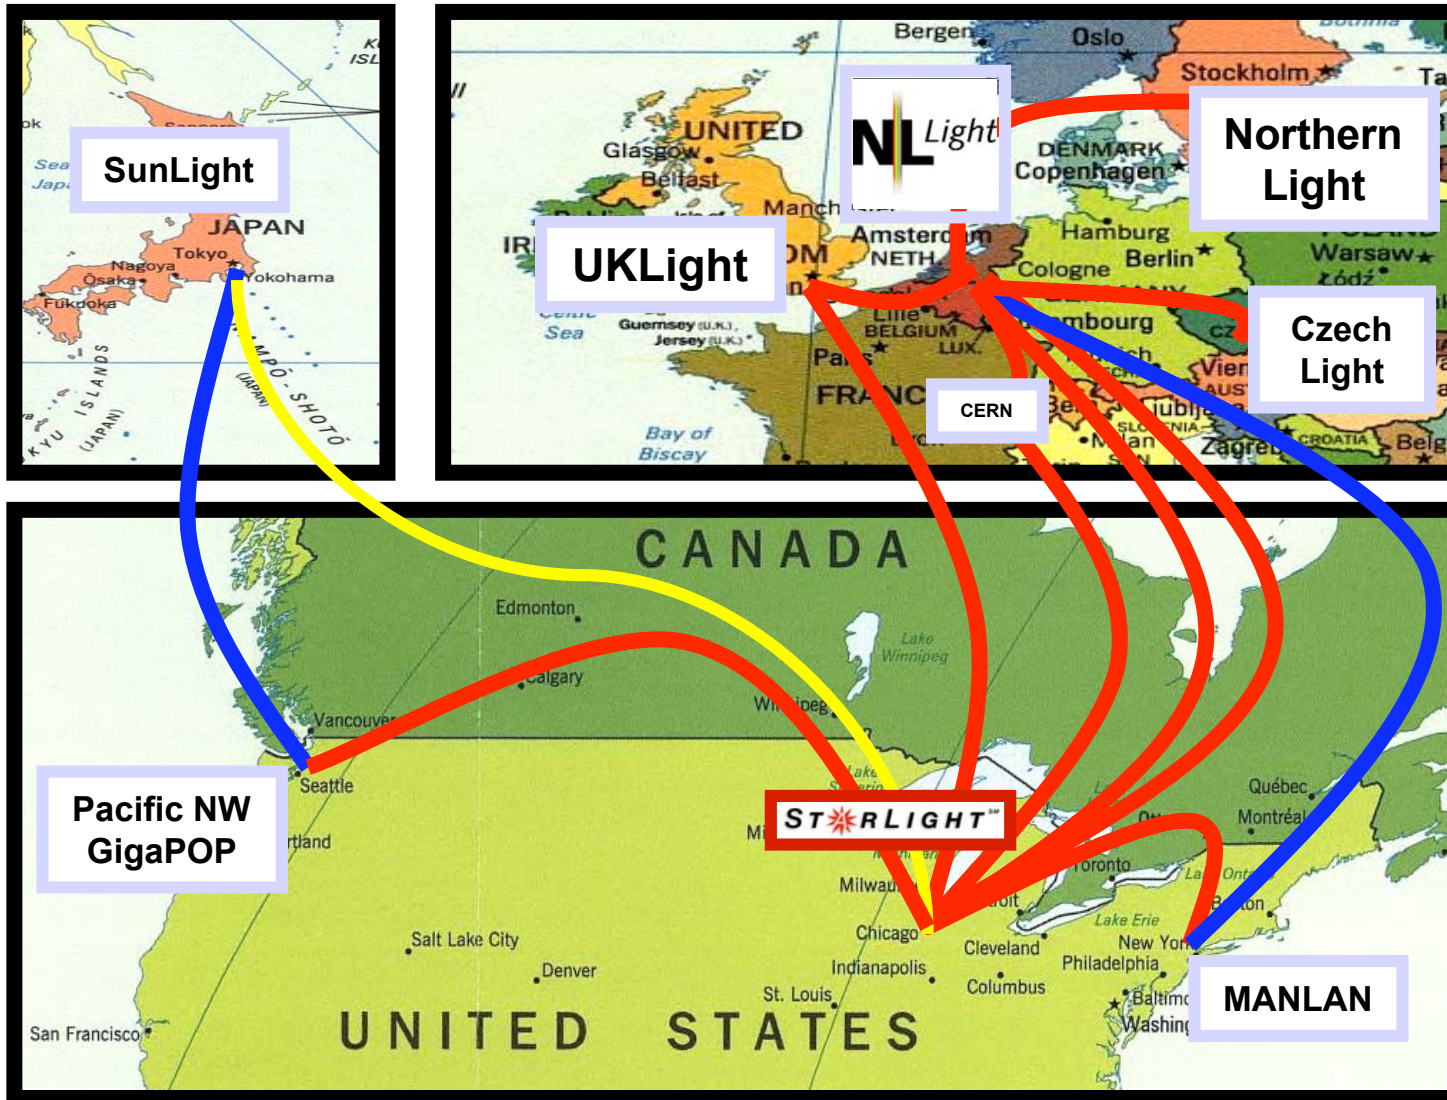
Source: Andrew Chien, UCSD
OptIPuter Software Systems Architect

The OptIPuter Will Become a National-Scale Collaboratory in 2004



Source: Tom West, CEO, NLR

An International-Scale OptIPuter is Operational over the First Set of 76 International GE TransLight Lambdas



European lambdas to US

—8 GEs Amsterdam—Chicago

—8 GEs London—Chicago

Canadian lambdas to US

—8 GEs Chicago—Canada

—NYC

—8 GEs Chicago—Canada

—Seattle

US lambdas to Europe

—4 GEs

Chicago—Amsterdam

—3 GEs Chicago—CERN

CERN

European lambdas

—8 GEs Amsterdam—CERN

—2 GEs

Prague—Amsterdam

—2 GEs

Stockholm—Amsterdam

—8 GEs

London—Amsterdam

TransPAC lambda

—1 GE Chicago—Tokyo

IEEAF lambdas (blue)

—8 GEs NYC—Amsterdam

—8 GEs Seattle—Tokyo

SIO's High Resolution OptIPuter Visualization Systems

Earth Sciences are an OptIPuter Driver

